

Distributed Sensing for Quickest Change Detection of Point Radiation Sources

Gene T. Whipps^{*†} Emre Ertin[†] Randolph L. Moses[†]

[†]The Ohio State University, ECE Department, Columbus, OH 43210

^{*}U.S. Army Research Laboratory, Adelphi, MD 20783

Abstract—We consider the problem of distributed detection of a radioactive source using a network of emission count sensors. Sensor nodes observe their environment and a central fusion node attempts to detect a change in the joint probability distribution due to the appearance of a hazardous source at an unknown time and location. We consider a minimax-type distributed change-point detection problem that minimizes detection delay for a desired false alarm rate. A statistical model of the radiation source detection problem is formulated where sensors observations are correlated with non-identical distributions. We first derive a centralized detection algorithm that is asymptotically optimal for vanishing false alarm rate. Then we analyze the performance loss, as measured by the detection latency, when sensor counts are quantized at each sensor node. The detection latency of the centralized rule provides a lower bound on performance for the proposed distributed method. The empirical results indicate that the distributed detection strategy provides a reasonable tradeoff between latency and information bandwidth.

I. INTRODUCTION

Detection of radioactive sources is a key component in the prevention of illicit transport and use of hazardous radioactive material. The detection and identification of radioactive point sources, particularly from a network of sensors, has received greater attention in the last decade. Legacy systems consist of a single large sensor node, such as a Radiation Portal Monitor, that is best operated at controlled access points and typically requires a specific operational protocol such as driving the source through a portal. More recent designs employ a network of smaller, less expensive sensor nodes that can monitor larger areas. Each type of system has its merits and limitations.

In this work, we consider distributed detection of a radioactive source from a network of sensors. Each sensor node measures an emission count over a sampling time interval, performs local processing, and transmits either raw count measurements or a processed local count statistic to a fusion center. Radioactive source detection systems typically rely on sensor nodes that report radioactive emission counts over time intervals on the order of seconds [1], [2]. The fusion center combines sensor node information to make a detection decision. The decision is to declare the null hypothesis in which no radiation source is present and the sensor counts represent a background emission level, versus the alternate hypothesis in which a radiation source is present.

In this paper, we consider an architecture in which each sensor node makes a local binary decision based on current

observations only, binary decisions are transmitted to a fusion center, and the fusion center executes a change-point rule. Additionally, local sensors need only transmit positive decisions, in effect censoring their transmissions so as to reduce network bandwidth requirements. We consider a time-evolutionary model in which at some unknown time γ a source appears, and the goal is to detect the presence of the source with minimal expected detection delay subject to a false alarm constraint.

Prior works on radiation detection have also considered detection methods where local binary decisions are made at each sensor node followed by a global decision rule at a central node. The present work differs from these works in two distinct ways. In prior works, the local decision rules are either batch [3] or sequential [2], [4], [5], and the global decision rule is a batch-type rule. In these problem formulations, the time when the radiation source appears is assumed to correspond to the first sample if the source is present. This type of formulation is well-matched to some operational protocols; an example case might be at a controlled checkpoint where vehicles are stopped and scanned. In contrast, we consider the case where the radiation source appears at an unknown point in time during continuous operation. Instead of batch or sequential tests, we apply quickest change-point tests at the fusion center. Specifically, we consider sensors making periodic observations, whose statistical distribution undergoes a change as a result of a point source appearing with unknown location and intensity. The objective is to detect the appearance of the point source as quickly as possible, subject to a constraint on the false alarm rate. In the literature, minimax and Bayesian rules are derived for single sensor source cases depending on whether the change is modeled as deterministic or random. Here we consider a network of geographically distributed sensors linked to a fusion center through band limited links and model the change point as unknown deterministic and the source parameters random with known prior distribution. Distributed quickest change detection methods [6] consider the case where information bandwidth between sensors and the fusion center is limited. A potential strategy for reducing bandwidth is to quantize observations at each sensor and communicate quantized values to the fusion center. Alternatively, sensors can perform localized change detection tests and communicate their local decisions to the fusion center. Optimal test procedures for these cases were

derived under the assumption that under each hypothesis the observations are independent. In our model of our radiation source detection problem sensor readings are correlated through the source location and intensity and therefore these results are not directly applicable. Recently, Qian, *et al.* [7] considered quickest change-point detection using a sensor network. They consider non-parametric CUSUM tests at each sensor node without an explicit statistical model of the source and propose a batch-type global decision rule at the central node. The proposed centralized and distributed tests are evaluated empirically through simulations.

In this paper, we use an explicit statistical model in which a radioactive source is characterized by an unknown feature vector characterizing its location and source strength, so sensor measurements are correlated with one another under the alternate hypothesis. When the post-change distribution is unknown and modeled as member of parametric family, one can follow a generalized likelihood ratio based approach [8] or a marginal likelihood ratio approach [9]. Here we consider the case of identical quantization rules at each sensor node and use a multi-sensor extension of the quickest detection rule with nuisance variables derived in [9] consistent with the correlated-observations model.

We provide a statistical model of the radiation source detection problem using sensors with correlated measurements with non-identical statistics and provide a centralized detection algorithm that is asymptotically optimal for vanishing false alarm rate. Then we analyze the performance loss, as measured by the detection latency, when sensor counts are quantized at each sensor node. The detection latency of the centralized rule provides a lower bound on detection latency for the proposed distributed method. Our empirical results show that while the achieved latency of the decentralized algorithm with binary quantization at the sensors is worse than the centralized scheme, except asymptotically small false alarm rates, it provides a reasonable tradeoff between latency and information bandwidth.

While the paper considers radioactive source detection as a motivating application, it is worth noting that the results of this paper apply to hypothesis testing in which the observations are count statistics whose count distributions are governed by Poisson distributions under the null and alternate hypotheses. As mentioned previously, the sensor node measurements are correlated under the alternate hypothesis case.

The remainder of this paper is outlined as follows. In Section II, we outline the sensor network model and derive a centralized quickest change-point decision rule. In Section III, we derive a distributed decision rule that reduces the communications overhead compared to the centralized rule. In Section IV, we present a numerical example validating analytic results. The example also demonstrates the increase in detection latency of the distributed rule compared to the centralized one. Finally, conclusions are given in Section V.

II. CENTRALIZED DECISION RULE

A. Problem Formulation

We consider a network of M sensor nodes, with index i denoting the i^{th} such node. Each node measures an emission count over an observation time period; we denote by $x_{i,1}, x_{i,2}, \dots, x_{i,n}$ the (random) time sequence of observed counts at the i^{th} sensor node.

For radiation source detection, we consider a binary change detection problem in which one wishes to determine whether a radioactive source has become present or not. The emission counts are random variables whose probability mass distribution changes at some unknown point in time γ , called a change point; note that γ an integer. We assume the change point occurs simultaneously at all sensor nodes, thus neglecting signal propagation delays. For nuclear radiation, the observation period, which is on the order of seconds [1], [2], is much longer than particle travel times, which are on the order of microseconds for a kilometer baseline, so the assumption of simultaneous change point times is reasonable for this detection application.

Prior to the change point, the emission counts $x_{i,n}$ for $n < \gamma$ are mutually independent across the sensor network with probability mass function p_i^0 at node i and where the 0 superscript denotes the null hypothesis H_0 . Note that the null hypothesis PMFs need not be identical. At time $n = \gamma$, we assume a source becomes present in the scene and is parameterized by an unknown parameter vector θ ; for example, $\theta = (\alpha, \rho, \phi)$ might characterize is count strength α and polar location (ρ, ϕ) , although other parameterizations may also exist. After the change point, the counts $x_{i,n}$ for $n \geq \gamma$ are *conditionally* mutually independent with probability mass function p_i^θ when conditioned on the unknown source parameters θ . In other words, observation $x_{i,\gamma}$ at node i is the first observation governed by the conditional PMF p_i^θ . The source parameter vector θ is assumed to be random with known prior F_θ . We note that, in contrast to [2] where all n observations are drawn from just one of two hypotheses, here a change in distribution can occur anywhere within the n -length time observation sequence.

In this paper, we assume that both the pre-change and post-change emission counts are Poisson distributed with intensities

$$\lambda_i^0 = b_i, \quad i = 1, 2, \dots, M, \quad (1)$$

under the pre-change case and

$$\lambda_i^\theta = a_i(\theta) + b_i, \quad i = 1, 2, \dots, M, \quad (2)$$

under the post-change case. The source intensity is modeled by $a_i(\theta) = \alpha(d_i + 1)^{-2}$, where d_i is the Euclidean distance between the i^{th} node and the source.¹

¹In [2], the emission counts at a sensor node due to the source have intensity α/d^2 . The denominator is modified here from d^2 to $(d+1)^2$ to avoid an unbounded intensity as the source-to-sensor distance approaches zero.

B. Quickest Change Detection

The quickest change detection problem can be formulated as follows. We consider a sequential detection problem in which at each point n , a decision rule processes the sequence of measurements $x_{i,k}$ for $1 \leq i \leq M$ and $1 \leq k \leq n$, and either declares that no change has occurred or declares that a change has occurred, and therefore that a radioactive source has become present at some point in time up to n . Let N denote the ‘stopping time’, which is the time at which a sequential decision rule declares that a change in distribution has occurred. The stopping time is random because it is a function of the sequence of emission counts. It is considered a false alarm if the rule declares a change point but no change in distribution actually occurred.

The design goal adopted here is to detect a change in distribution with the least amount of delay and for a given false alarm constraint. The probability of false alarm cannot be modeled analytically because γ or its distribution are unknown. Thus, in a non-Bayesian setting, when no prior is assumed on the change point γ , the constraint is on the false alarm rate, which is typically defined by the average time between false alarms. Specifically, the user selects the minimum average time between false detections, denoted by A , and the false alarm constraint is defined as $E_0(N) \geq A$. A minimax-type criterion on the detection delay (i.e., stopping time) was proposed in [10] and is defined by

$$DD(N) = \sup_{n \geq 1} E_n(N - n | N \geq n). \quad (3)$$

The problem is then to find the decision rule that minimizes the detection delay $DD(N)$ such that $E_0(N) \geq A$.

C. Centralized Quickest Change Detection

We first consider the centralized case, in which there is $M = 1$ sensor node. We further assume that the pre-change and post-change distributions— p^0 and p^θ respectively—are within the class of exponential distributions. Under this assumption, it was shown in [9] that the composite rule given by

$$N = \inf \left\{ n : \max_{1 \leq k \leq n} \int_{\Theta} \prod_{j=k}^n \frac{p^\theta(x_{1,j})}{p^0(x_{1,j})} dF_\theta \geq A \right\}, \quad (4)$$

is asymptotically optimal under the false alarm constraint if the following holds:

$$DD^\theta(N) \sim \frac{|\log A|}{D(p^\theta || p^0)} \quad \text{as } A \rightarrow \infty. \quad (5)$$

provided $D(p^\theta || p^0) > 0$ for every θ in the support of F_θ . Here, the notation $g(x) \sim h(x)$ as $x \rightarrow x_0$ is used to denote that $g(x) = h(x)(1 + o(1))$. Also, the notation DD^θ means that the expectation in equation (3) is replaced by a conditional expectation with conditioning on the particular θ . The value $D(p^\theta || p^0)$ in equation (5) is the relative entropy between distributions p^θ and p^0 .

We note that for the current problem formulation, both the pre-change and post-change conditional distributions are exponential distributions of the form $f_\lambda(x) =$

$h(x) \exp\{\eta(\lambda)T(x) - A(\eta)\}$ with $h(x) = 1/x!$, $\eta(\lambda) = \log \lambda$, $A(\eta) = e^{\eta(\lambda)}$, and $T(x) = x$.

We extend the single-node decision rule (4) to the multi-node case with heterogeneous Poisson statistics when conditioned on the source parameters. The single-node distributions are replaced by the multi-node joint distributions, and the likelihood ratio, conditioned on the source parameters, at time j becomes

$$\frac{p^\theta(x_{1,j}, x_{2,j}, \dots, x_{M,j})}{p^0(x_{1,j}, x_{2,j}, \dots, x_{M,j})} = \prod_{i=1}^M \frac{p_i^\theta(x_{i,j})}{p_i^0(x_{i,j})}.$$

The decision rule (4) is expressed in [9] as a function of the accumulation of the observations. Similarly, if we define the cumulative counts as $s_{i,n} = \sum_{j=1}^n x_{i,j}$ with $s_{i,0} = 0$. The centralized decision rule is then given by

$$N_c = \inf \left\{ n : \max_{1 \leq k \leq n} \int_{\Theta} \prod_{i=1}^M f_i^{n-k}(s_{i,n} - s_{i,k-1}; \theta) dF_\theta \geq A \right\} \quad (6)$$

where

$$f_i^m(x; \theta) = \exp \left(x \log \left(\frac{\lambda_i^\theta}{\lambda_i^0} \right) - (m+1)(\lambda_i^\theta - \lambda_i^0) \right).$$

Unfortunately, this rule does not appear to have a recursive form. As a result, (6) requires the storage of the entire history of the cumulative counts $s_{i,1}, s_{i,2}, \dots, s_{i,n}$ for each sensor node in the network. Furthermore, the computational complexity increases proportionally with n ; the rule is more cumbersome to compute with each new observation. Sliding window rules are proposed in [11] to constrain the memory and computation burdens and are shown to demonstrate comparable performance. In our simulations we have used a sliding window rule to limit the computational complexity.

Being conditionally independent and Poisson distributed across the network, the relative entropy between the emission count distributions is

$$D(p^\theta || p^0) = \sum_{i=1}^M \lambda_i^\theta \log \left(\frac{\lambda_i^\theta}{\lambda_i^0} \right) - (\lambda_i^\theta - \lambda_i^0), \quad (7)$$

given the source parameter θ . Equations (7) and (5) provide a simple metric for evaluating the detection latency of the centralized decision rule for feasible sources.

III. DISTRIBUTED DECISION RULE

The centralized detection rule considers the case where Poisson counts are periodically reported to the fusion center without quantization at regular intervals. This could be a non-viable method when the communication bandwidth allocated to the sensor network is constrained. Under this scenario, the sensor nodes can employ two potential strategies to reduce the information bandwidth of their messages. In a distributed decision structure sensors can run local tests to detect the change point and communicate their detections to the fusion center sporadically. It is very hard to derive optimal distributed

decision strategies for this setup with no feedback from the fusion center due to the non-classical information structure of the problem [6], where each decision maker is operating with different information history or filtration structure. In addition, for practical scenarios when the links are unreliable, this decision structure could be very brittle as the loss of one local detection message can result in unacceptable false-negatives.

Alternatively, sensors can quantize their counts and communicate a quantization index to the fusion center periodically. This strategy is inherently more robust to message losses as messages indicating high count quantization level will be repeatedly transmitted after the appearance of the source. This could be considered as a special case of the decision rule in (6) with the sensor observation statistics replaced by sensor message statistics. In this paper, we consider this strategy and investigate the choice of quantization thresholds to optimize system level performance. Here we choose binary quantizers where counts are directly thresholded with a local threshold τ , resulting in sensor messages that are binomially distributed conditioned on sensor parameter vector θ .

For the mean-shift Poisson model described in Section II, it is straightforward to show the likelihood ratios are monotone non-decreasing as a function of an emission count. Thus, binary quantization via simple thresholding of the likelihood ratio of a measured count is equivalent to thresholding a measured count directly. Each sensor node then transmits a binary value corresponding to a measured emission count with conditional probability $q_i^0(\tau) = \Pr(x_{i,j} > \tau)$, under the pre-change case and $q_i^\theta(\tau) = \Pr(x_{i,j} > \tau; \theta)$, under the post-change case. We simplify the notation to $q_i^0 = q_i^0(\tau)$ and $q_i^\theta = q_i^\theta(\tau)$ with the understanding that the probabilities depend on threshold τ . These probabilities are points, determined by τ , on the complementary cumulative distribution function of Poisson variables with intensities λ_i^0 and λ_i^θ , respectively. The distributed decision rule is a function of the accumulation of binary values. Define $u_{i,j}$ as the quantized value of the emission count $x_{i,j}$, and $s_{i,n}^d = \sum_{j=1}^n u_{i,j}$ with $s_{i,0}^d = 0$. The distributed change-point rule is then given by

$$N_d = \inf \left\{ n : \max_{1 \leq k \leq n} \int_{\Theta} \prod_{i=1}^M g_i^{n-k}(s_{i,n}^d - s_{i,k-1}^d; \theta) dF_\theta \geq A \right\} \quad (8)$$

where

$$g_i^m(x; \theta) = \exp \left(x \log \left(\frac{q_i^\theta}{q_i^0} \right) - (m+1-x) \log \left(\frac{1-q_i^\theta}{1-q_i^0} \right) \right).$$

To evaluate the detection latency, the relative entropy in equation (5) is replaced by the relative entropy between the Bernoulli distributions of the quantized counts, which is simply

$$D(p^\theta || p^0) = \sum_{i=1}^M q_i^\theta \log \left(\frac{q_i^\theta}{q_i^0} \right) + (1 - q_i^\theta) \log \left(\frac{1 - q_i^\theta}{1 - q_i^0} \right). \quad (9)$$

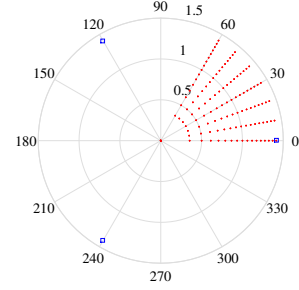


Fig. 1. Sensor locations are marked by blue squares and possible source locations are marked by red dots. Sensors are uniformly spaced on the circumference of a circle with radius $r_0 = \sqrt{2}$ units.

Using equations (5), (7), and (9), comparisons of the detection latency for low false alarm rates can be made between the centralized and decentralized detectors. The detection latency of both the centralized and distributed decision rules are dependent on the signal-to-noise ratio and the decision threshold A . While a larger threshold A results in a lower false alarm rate, the average detection latency increases. As one would expect and as seen through the relative entropies, decreasing signal-to-noise ratios surely increases on average the detection latency. While not shown here explicitly, an application of Jensen's inequality would show that the relative entropy of the distributed method is no greater than that of the centralized method. As a result, we would expect the average detection latency of the centralized detector to be a lower bound to that of the distributed detector. Unlike the centralized case, the detection latency of the distributed decision rule also depends on the quantization threshold τ .

IV. NUMERICAL SIMULATION

We present a simulation example to evaluate the efficacy of the distributed change point rule given in (8), where sensor counts are quantized to binary statistics, for detecting the presence of a radiation source and compare it to the performance of the centralized rule given in (6). The radiation source is parameterized by a vector $\theta = (\alpha, \rho, \phi)$ where (ρ, ϕ) denotes polar location and α denotes source intensity. We model the source location as uniformly distributed within a circular region \mathcal{R} centered at $(0, 0)$ and with radius $r_0 = \sqrt{2}$ units. The source intensity is modeled as uniformly distributed in an interval $[a_1, a_2]$.

We consider a sensor network with $M = 3$ nodes, located on the circumference of \mathcal{R} at polar angles $0, 2\pi/3$, and $4\pi/3$ (see Figure 1). We assume that the ambient radiation count levels under the null hypothesis are Poisson with levels given by $b_1 = b_2 = b_3 = 10$ emission units as in equation (1). We note that even if these ambient levels are unknown, they can be recursively estimated with the estimation error made negligible under practical scenarios. When a source is present, the sensors measure a radiation count whose Poisson levels increase by $a_i(\theta) = \alpha(d_i + 1)^{-2}$ as in equation (2).

In the simulations that follow, we compute empirical averages by locating the radiation source at a discrete set of loca-

tions as shown in Figure 1, with 7 polar angle discretizations in $[0, \pi/3]$ and 17 radial distance discretizations in $[0, \sqrt{2}]$. Due to the symmetry of the problem, this configuration is equivalent to covering the entire region \mathcal{R} . The source intensities are uniformly distributed in the interval $[10, 20]$ emission units for equivalent signal-to-noise ratios (SNRs) in the range of 0 dB–6 dB. Figure 1 plots the locations of the sensor nodes (blue squares) and feasible source locations (red dots).

The hypotheses are simulated through Monte-Carlo trials with change point $\gamma = \infty$ (H_0) and $\gamma = 1$ (H_1), respectively. Each trial runs until the decision rule declares a detection. We consider expected time between false alarms of $A = 10$ to $A = 10^6$. Under H_0 , each trial run may be quite long because it is desired to have $E_0(N) \geq A$. Thus, we limit the largest value of A to be 100 for the H_0 trials in order to compute sample statistics in a reasonable amount of time.

Figure 2 shows the empirical false alarm rates of the centralized and distributed detectors versus A , for 10^3 Monte-Carlo trials. The empirical false alarm rates are computed as the inverse of the sample average of stopping times N_c and N_d for the centralized and distributed detectors, respectively. The false alarm rates for the distributed detector are evaluated for quantization thresholds $\tau = 10$, $\tau = 12$, and $\tau = 14$. As seen in the figure, the false alarm rates of the distributed method are very comparable with that of the centralized one. The empirical rates also appear inversely proportional to A as desired. Curiously, the false alarm rates of both detectors are nearly an order of magnitude less than the desired maximum.

Figure 3 shows the empirical average detection latency versus the decision threshold A and for the three different quantization thresholds τ . Also shown is the latency lower bound as computed empirically from trials using the centralized detector. It can be seen that, even for a small number of $M = 3$ sensors, modest loss of detection delay is realized across a wide range of false alarm rates. Figure 4 shows how latency changes as a function of the quantization threshold for several user-selected false alarm rates (as governed by threshold A). For this configuration of the sensor network and source distribution, the threshold $\tau = 12$ results in the lowest detection latency for the distributed detector. These results are obtained by averaging across the range of θ realizations for location and amplitude from 10^4 Monte-Carlo trials.

To understand the asymptotic optimality, for vanishing false alarm rates, of the change-point detectors we consider three extreme cases, one for each corner of the wedge (see in Figure 1) of feasible source locations. We conjecture that detection delay for other cases will fall somewhere between these extremes. First, Figure 5 shows latency performance for a source located at $(\rho, \phi) = (\sqrt{2}, 0)$ (that is, co-located with one of the sensors) and intensity $\alpha = 10$. Both asymptotic analytic latency performance (dashed curves) and empirical detection latency (solid curves) from 10^4 Monte-Carlo trials are shown, for both the centralized detection algorithm and its distributed counterpart. We see very good agreement between analytical and empirical latency across a range of thresholds. The slope of empirical detection delay very nearly matches

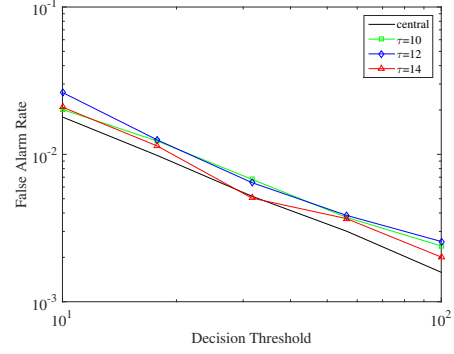


Fig. 2. Empirical false alarm rate from 10^3 Monte-Carlo trials versus the decision threshold A for the centralized and distributed methods.

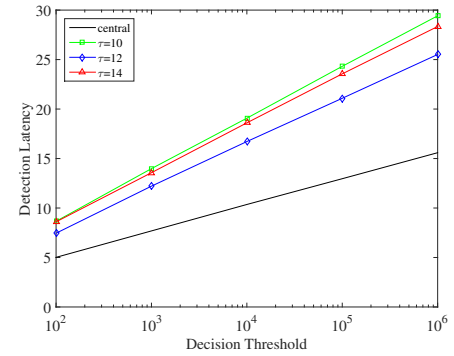


Fig. 3. Empirical average detection latency from 10^4 Monte-Carlo trials versus the decision threshold A for the centralized and distributed methods.

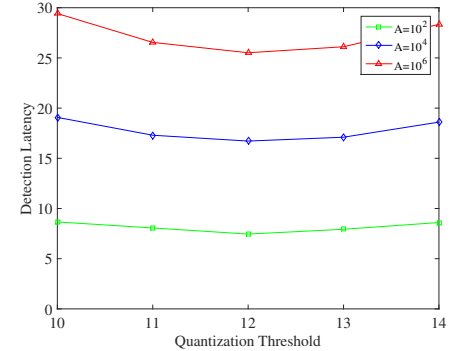


Fig. 4. Empirical average detection latency from 10^4 Monte-Carlo trials versus the quantization threshold τ for the distributed method.

the desired asymptotic slope to within a small value that does not appear to depend on A .

Figures 6 and 7 show latency performance for a source located at $(\rho, \phi) = (0, 0)$ (that is, maximally far from all three sensors) and $(\rho, \phi) = (\sqrt{2}, \pi/3)$. The intensities are $\alpha = 10$ and $\alpha = 20$, respectively. Here again we see good agreement between asymptotic analytic predictions and empirical results. In each case, the centralized detection rule outperforms the distributed system with quantized counts with roughly 50% reduction in the detection latency. However, when information

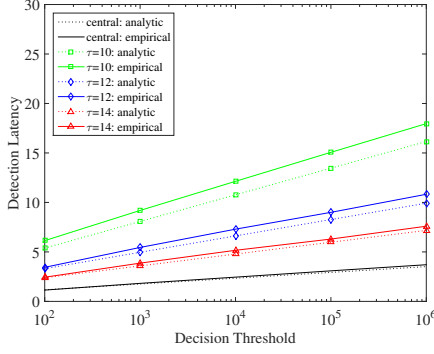


Fig. 5. Average latency versus the decision threshold A for a source with parameters $\theta = (10, \sqrt{2}, 0)$.

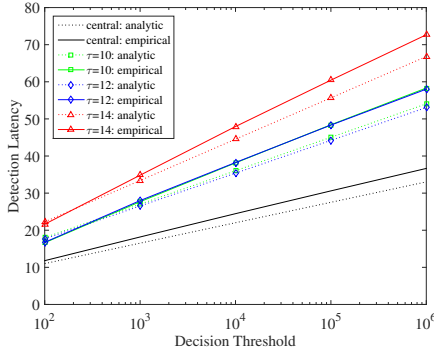


Fig. 6. Average latency versus the decision threshold A for a source with parameters $\theta = (10, 0, 0)$.

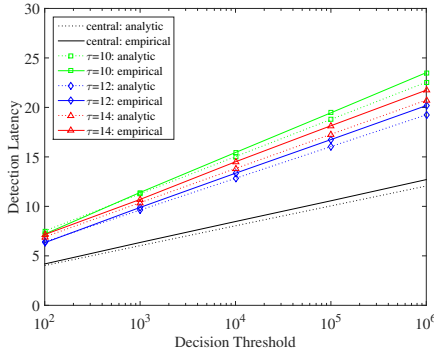


Fig. 7. Average latency versus the decision threshold A for a source with parameters $\theta = (20, \sqrt{2}, \pi/3)$.

bandwidth is constrained by the communications network, the centralized detector may not be practical.

V. CONCLUSIONS

We have considered the problem of detecting the onset of a radiation source from a network of sensor nodes, when the onset time is unknown. We adopted a quickest change detection formulation, in which the design goal is to minimize the delay between the onset time and the average detection time, subject to a false alarm rate constraint. The problem

formulation differs from earlier ones in that the onset time is unknown, and the sensor node measurements under the source-present hypothesis are no longer independent due to the correlated nature of count rates that result from correlated distances to the radiation source. The structure of the distributed decision rule is one in which each sensor node performs a local detection test and transmits only positive detections at each time slot; by transmitting only positive detections, nodes in effect self-censor, and network bandwidth is efficiently used. We presented numerical examples and showed both modest performance loss with respect to a fully centralized detector and good agreement with asymptotic performance analyses over a range of source parameter values. While the results of the paper have focused on detection of radiation sources, the detection architecture applies to other applications in which Poisson count statistics are collected by a network of sensors. Here, we considered a simple line search to search for the binary quantization threshold. This direct optimization method will be intractable for the case of multi-level quantizers. One potential strategy is to use a proxy, such as maximizing the minimum Kullback–Leibler divergence as taken over the support of the source parameters.

ACKNOWLEDGMENT

The research reported here was partially supported by the U.S. Army Research Laboratory and by a grant from ARO, W911NF-11-1-0391.

REFERENCES

- [1] S. M. Brennan, A. M. Mielke, D. C. Torney, and A. B. MacCabe, "Radiation detection with distributed sensor networks," *Computer*, vol. 37, no. 8, pp. 57–59, 2004.
- [2] N. S. Rao, M. Shankar, J.-C. Chin, D. K. Yau, S. Srivathsan, S. S. Iyengar, Y. Yang, and J. C. Hou, "Identification of low-level point radiation sources using a sensor network," in *Information Processing in Sensor Networks, 2008. IPSN'08. International Conference on*. IEEE, Apr. 2008, pp. 493–504.
- [3] A. H. Liu, J. J. Bunn, and K. M. Chandy, "An analysis of data fusion for radiation detection and localization," in *FUSION*, 2010, pp. 1–8.
- [4] N. S. Rao, C. W. Glover, M. Shankar, J.-C. Chin, D. K. Yau, C. Y. Ma, Y. Yang, and S. Sahni, "Improved spst detection using localization with application to radiation sources," in *Information Fusion, 2009. FUSION'09. 12th International Conference on*. IEEE, 2009, pp. 633–640.
- [5] K. M. Chandy, J. Bunn, and A. Liu, "Models and algorithms for radiation detection," in *Modeling and Simulation Workshop for Homeland Security*, 2010, pp. 1–6.
- [6] V. V. Veeravalli, "Decentralized quickest change detection," *Information Theory, IEEE Transactions on*, vol. 47, no. 4, pp. 1657–1665, 2001.
- [7] L. Qian, J. Fuller *et al.*, "Quickest detection of nuclear radiation using a sensor network," in *Homeland Security (HST), 2012 IEEE Conference on Technologies for*. IEEE, 2012, pp. 648–653.
- [8] G. Lorden, "Procedures for reacting to a change in distribution," *The Annals of Mathematical Statistics*, vol. 42, no. 6, pp. 1897–1908, Dec. 1971.
- [9] M. Pollak and D. Siegmund, "Approximations to the expected sample size of certain sequential tests," *The Annals of Statistics*, vol. 3, no. 6, pp. 1267–1282, Nov. 1975.
- [10] M. Pollak, "Optimal detection of a change in distribution," *The Annals of Statistics*, vol. 13, no. 1, pp. 206–227, Mar. 1985.
- [11] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Transactions on Information Theory*, vol. 44, no. 7, pp. 2917–2929, Nov. 1998.